WEST

**End of Result Set**

Generate Collection | Print

L1: Entry 1 of 1 File: TDBD Sep 1, 1992

TDB-ACC-NO: NA9209404

DISCLOSURE TITLE: Limited Distributed DASD Checksum, a RAID Hybrid.

PUBLICATION-DATA:
IBM Technical Disclosure Bulletin, September 1992, US

VOLUME NUMBER: 35
ISSUE NUMBER: 4A
PAGE NUMBER: 404 - 405

PUBLICATION-DATE: September 1, 1992 (19920901)

CROSS REFERENCE: 0018-8689-35-4A-404

DISCLOSURE TEXT:

- A Redundant Array of Inexpensive Disk (RAID) that fully distributes parity across all arms (RAID-5) performs better than a RAID that has a single parity drive (RAID-4). But with RAID-5, the usable capacity (viewed by the system) of all disk units changes when units are added or removed from the checksum set. This adds to the complexity of attaching a RAID-5 box to an existing system. - This article describes a hybrid solution that performs better than RAID-4 and allows growth without changing the usable capacity of the existing units in the checksum set. - First, some background information. A redundant array with inexpensive disk (RAID) provides protection from data loss in case of a single unit failure. Parity in the disk array is used to reconstruct data. This improves the availability of the data. - A redundant array with a single parity unit is called a RAID-4 disk drive (DASD) array. Data is kept on the N units while parity is kept on the N+1 unit. An example with N=3 is shown in Fig. 1. - The system has addressability to the full capacity of the N units but cannot address the parity unit. One advantage of the RAID-4 architecture is that adding or removing units is relatively simple because the change does not affect the other N units. The parity is simply regenerated for the N units in the set. The major disadvantage of RAID-4 is the P unit becomes a performance bottleneck during write operations. - A RAID-5 DASD array removes the performance bottleneck by distributing the parity across all units. The system has addressability to all of the units but each has reduced capacity. The example shown in Fig. 2 has four units, thus the system can address only 75% (3/4) of each unit. - The major disadvantage of the RAID-5 architecture is adding or removing units from the checksum set changes the percentage of user storage on all the units in the set. For example, if two units were added to the configuration shown in Fig. 2, parity would be redistributed across all six units. The system would have access to 83.3% (5/6) of each unit, a growth of 8.3% per unit. Similarly, if units are later removed, the remaining units reduce their user capacity. Handling a DASD unit that changes in size adds considerable complexities to the AS/400* system design. A unique model is required for each one of the possible checksum configurations in the box. As units are added or removed, system storage and data management must recognize model changes and handle the addressing changes appropriately. Customers that have multiple Auxiliary Storage Pools (ASPs) configured across the DASD box may find them unmanageable because adding or removing a unit would affect all of those pools. - A solution that improves the RAID-4 performance and simplifies the growth complexities inherent in RAID-5 is shown in Fig. 3. - The checksum is distributed yet limited to a subset of the units in the checksum set. The first four units in this example have the parity for all six units in the checksum set. Those units provide the system with a fraction of their full capacity (just like a RAID-5

architecture). The other units would be full capacity (just like a RAID-4 architecture). The major advantage of this solution is that the sizes of each unit does not change when units are added and removed outside the limited checksum set. Another advantage of this solution is the performance is better than the RAID-4 architecture. The solution also removes the need for many unique model numbers and simplifies ASP management for both the system and customer. The disadvantages are that the system must handle two different DASD sizes in a box and require a minimum number of units in the distributed checksum set. These are acceptable compromises if the goals are to allow growth of the RAID configuration and reduce the overall complexity of the box attachment to the system.